arXiv:math/0411071v2 [math.PR] 23 Apr 2005

A coalescent model for the effect of advantageous mutations on the genealogy of a population

by Rick Durrett* and Jason Schweinsberg†

February 1, 2008

Abstract

When an advantageous mutation occurs in a population, the favorable allele may spread to the entire population in a short time, an event known as a selective sweep. As a result, when we sample n individuals from a population and trace their ancestral lines backwards in time, many lineages may coalesce almost instantaneously at the time of a selective sweep. We show that as the population size goes to infinity, this process converges to a coalescent process called a coalescent with multiple collisions. A better approximation for finite populations can be obtained using a coalescent with simultaneous multiple collisions. We also show how these coalescent approximations can be used to get insight into how beneficial mutations affect the behavior of statistics that have been used to detect departures from the usual Kingman's coalescent.

1 Introduction

Our goal in this paper is to describe the coalescent processes that arise when we consider the genealogy of a population that is affected by repeated beneficial mutations. The starting point for this analysis will be the continuous-time population model introduced by Moran (1958). In this model, the population size is fixed at $2N$. Each individual independently lives for a time that is exponentially distributed with mean 1 and then is replaced by a new individual. The parent of the new individual is chosen at random from the $2N$ individuals, including the one being replaced. Note that we can think of the population as consisting of $2N$ chromosomes of N diploid individuals, so each member of the population has just one parent.

Suppose we sample n individuals at random from this population at time zero. To describe the genealogy of the sample, we will define the ancestral process, which will be a continuous-time Markov process $(\Psi_N(t), t \geq 0)$ whose state space is the set $\mathcal{P}_n$ of partitions of $\{1, \ldots, n\}$. The ancestral process describes the coalescence of lineages as we follow the ancestral lines of the sampled individuals backwards in time. More precisely, $\Psi_N(0)$ is the partition of $\{1, \ldots, n\}$ into n singletons, and $\Psi_N(t)$ is the partition of $\{1, \ldots, n\}$ such that i and j are in the same block of $\Psi_N(t)$ if and only if the ith and jth individuals in the sample have the same ancestor at time $-Nt$. It is well-known that the process $(\Psi_N(t), t \geq 0)$ is Kingman's coalescent, a coalescent

*Partially supported by NSF grants from the probability program (9877066 and 0202935) and from a joint DMS/NIGMS initiative to support research in mathematical biology (0201037).

†Supported by an NSF Postdoctoral Fellowship

process introduced by Kingman (1982). Kingman's coalescent is a $\mathcal{P}_n$-valued Markov process that starts from the partition of $\{1, \ldots, n\}$ into singletons. All transitions involve exactly two blocks of the partition merging together, and each such transition occurs at rate one.

Within the last decade, progress has been made on describing the genealogy of populations in models that allow for natural selection. Krone and Neuhauser (1997) and Neuhauser and Krone (1997) studied a model in which each individual can be of type 1 or 2. An individual of type i produces offspring at rate λ_i, with $\lambda_2 > \lambda_1$ so that type 2 is advantageous. Each new offspring replaces a randomly chosen individual from the population, and is the same type as its parent with probability $1 - u_N$ and the opposite type with probability u_N. Under certain assumptions, they show that the genealogy of a sample from the population can be described using what they call an ancestral selection graph. Additional work of Donnelly and Kurtz (1999) and Barton, Etheridge, and Sturm (2004) has incorporated recombination as well as selection into the model.

The ancestral selection graph arises in the limit as $N \to \infty$ in the case of weak selection, where the selective advantage $\lambda_2/\lambda_1 - 1$ and the mutation rates u_N are $O(1/N)$. Then, as $N \to \infty$ the fraction of individuals with the favored allele can be approximated by a diffusion process. In this paper, we consider strong selection, where the selective advantage is $O(1)$. With strong selection, when a beneficial mutation occurs, there is a positive probability that the beneficial allele will spread to the entire population, an event known as a selective sweep.

At the end of a selective sweep, the entire population has the favorable allele, and every member of the population will trace that favorable allele back to the individual that had the beneficial mutation that caused the selective sweep. However, the genealogy becomes more complicated when we consider recombination. Diploid individuals usually do not inherit an identical copy of one of their parent's chromosomes. Instead, the inherited chromosome consists of pieces of each of a parent's two chromosomes. Since a chromosome is coming from two places, we need to consider the genealogy not of an entire chromosome but of a particular site of interest on the chromosome. When a selective sweep is caused by a beneficial mutation at a site other than the site of interest, many individuals may trace their gene at the site of interest back to the individual that had the beneficial mutation at the beginning of the selective sweep, while others may trace their gene at the site of interest to a different ancestor because of recombination between the two sites on the chromosome. This effect was first studied by Maynard Smith and Haigh (1974), who called it the "hitchhiking effect."

As we will show, the typical duration of a selective sweep is only $O(\log N)$. Therefore, when we speed up time by a factor of N to define the ancestral process, the selective sweep takes place almost instantaneously. Consequently, if we sample n individuals some time after a selective sweep and define the ancestral process as before, the ancestral process behaves like Kingman's coalescent until we get back to the time of a selective sweep. At that time, many lineages may coalesce because they get traced back to the individual with the mutation that caused the selective sweep. This possibility was observed by Gillespie (2000), who referred to the resulting coalescent process as the "pseudohitchhiking model." We will show that if selective selective sweeps happen repeatedly throughout the history of a population at times of a Poisson process, as proposed by Gillespie (2000), then under suitable assumptions the ancestral processes will converge as $N \to \infty$ to a coalescent with multiple collisions, which is a $\mathcal{P}_n$-valued Markov process in which many blocks of the partition can merge at once into a single block. These coalescent processes were introduced by Pitman (1999) and Sagitov (1999).

While coalescents with multiple collisions are the limiting coalescent processes as $N \to \infty$,

an improved approximation for finite N can be obtained using a coalescent with simultaneous multiple collisions. Coalescents with simultaneous multiple collisions, which were introduced by Schweinsberg (2000) and Möhle and Sagitov (2001), are coalescent processes in which many blocks can merge at once into a single block, and many such mergers can occur simultaneously. They provide a better approximation than coalescents with multiple collisions in this context because, as noted by Barton (1998), Durrett and Schweinsberg (2004a), and Schweinsberg and Durrett (2004), multiple groups of lineages can coalesce at the time of a selective sweep.

Coalescents with multiple or simultaneous multiple collisions arise as limits of ancestral processes in populations that occasionally have very large families because ancestral lines that go back to an individual with many offspring will coalesce at the same time. Coalescents with multiple collisions arise when a single large family is possible in a given generation, while coalescents with simultaneous multiple collisions arise when one generation can contain many large families. For more details, see Sagitov (1999, 2003), Möhle and Sagitov (2001), and Schweinsberg (2003). The results in this paper provide a different biological application of these coalescent processes.

The rest of this paper is organized as follows. In section 2, we describe our model for how the population evolves when there can be beneficial mutations. We state our main result, which is that the genealogy of this process converges to a coalescent with multiple collisions. In section 3, we present the improved approximation involving a coalescent with simultaneous multiple collisions. The next two sections are devoted to applications of these results. In section 4, we discuss how multiple mergers affect the number of segregating sites and pairwise differences in a sample of DNA. These quantities are used in Tajima's D-statistic (see Tajima (1989)), which can be used to detect departures from the standard Kingman's coalescent. In section 5 we discuss how multiple mergers affect the number of mutations that appear on just a single individual in the sample, which is relevant to the test proposed by Fu and Li (1993) for detecting departures from Kingman's coalescent. Our results suggest that Fu and Li's test should have less power to detect selective sweeps, at least in large samples, than Tajima's D-statistic. Finally, in section 6, we prove the convergence and approximation theorems stated in sections 2 and 3.

2 Convergence to a coalescent with multiple collisions

In this section, we give a precise description of our model of a population that experiences beneficial mutations, and we state our main convergence theorem. We describe what happens following a single beneficial mutation in subsection 2.1, and we consider recurrent beneficial mutations in subsection 2.2. Then in subsection 2.3, we state the convergence result and give some examples.

2.1 The effect of a single beneficial mutation

In this subsection we describe how the population evolves after one of the $2N$ individuals experiences a beneficial mutation. We will denote the new favorable allele by B and the other allele by b. We assume the relative fitnesses of the two alleles are 1 and $1-s$, so the B alleles will tend to survive longer. Immediately after the mutation, one individual has the B allele and $2N-1$ have the b allele. Kaplan, Hudson, and Langley (1989) and Stephen, Wiehe, and Lenz (1992) proposed modeling the fraction of individuals $p(t)$ with the B allele at time t by using the logistic

differential equation

$$\frac{dp}{dt} = sp(1-p).$$

This approach has been popular in simulation studies. However, Durrett and Schweinsberg (2004a) showed that this approximation is not very accurate. Consequently, we will consider instead a modification to the Moran model that was studied by Durrett and Schweinsberg (2004a) and Schweinsberg and Durrett (2004).

At one site, each chromosome has a B or b allele, but we will be interested in the genealogy at another neutral site at which all alleles have the same fitness. As in the Moran model, each individual survives for a time that is exponentially distributed with mean 1, and then a replacement is proposed in which the parent of the proposed new individual is chosen at random from the $2N$ members of the population. However, to account for natural selection, whenever a replacement of a B chromosome with a b chromosome is proposed, the change is rejected with probability s. Also, to incorporate recombination into the model, we say that when a new individual is born, it inherits its alleles at both sites from the same parent with probability $1-r$. However, with probability r, there is recombination between the two sites, so the new individual inherits its allele at the neutral site from its parent's other chromosome. Because we are treating an individual's two chromosomes as two separate members of the population, we model this by saying that, with probability r, the new individual inherits the two alleles from two ancestors chosen independently at random from the population.

Suppose the beneficial mutation appears on one chromosome at time 0, and let $X(t)$ be the number of chromosomes with the favorable allele at time t. Let $\tau = \inf\{t : X(t) \in \{0, 2N\}\}$ be the time at which either the B or b allele disappears from the population. Suppose we take a random sample of n individuals from the population at time τ. Let Θ be the partition of $\{1, \ldots, n\}$ such that i and j are in the same block of Θ if and only if the ith and jth individuals in the sample have the same ancestor at time zero when we follow the ancestral lines associated with the neutral site of interest. The partition Θ then describes how the beneficial mutation affects the genealogy of the sample. We have the following result concerning the distribution of Θ. Here $Q_{p,n}$, for $p \in [0,1]$, is the distribution of a random partition Π obtained as follows. First, define a sequence of independent random variables $(\xi_i)_{i=1}^n$ such that $P(\xi_i = 1) = p$ and $P(\xi_i = 0) = 1-p$ for $i = 1, \ldots, n$. Then define Π such that one block of Π consists of $\{i \le n : \xi_i = 1\}$ and the remaining blocks of Π are singletons.

Proposition 2.1. *Fix $n \in \mathbb{N}$, and fix $s \in (0,1)$. Assume there is a constant C' such that $r \le C'/(\log N)$ for all N. Let $\alpha = r\log(2N)/s$, and let $p = e^{-\alpha}$.*

1. *There exists a positive constant C, depending continuously on s and α but not depending on N, such that $|P(\Theta = \pi | X(\tau) = 2N) - Q_{p,n}(\pi)| \le C/(\log N)$ for all $\pi \in \mathcal{P}_n$.*

2. *Let κ_0 be the partition of $\{1, \ldots, n\}$ into singletons. There exists a constant C, depending continuously on s and α but not depending on N, such that $P(\Theta \ne \kappa_0 \text{ and } X(\tau) = 0) \le CN^{-1/2}$.*

Note that in this proposition, the selective advantage s is assumed to be fixed, but the recombination probability r depends on N. Part 1 of the proposition, which is a restatement of Theorem 1.1 of Schweinsberg and Durrett (2004), implies that as $N \to \infty$, the distribution of Θ,

conditional on the event that a selective sweep occurs, converges to $Q_{p,n}$, where p represents the approximate fraction of lineages that coalesce at the time of the selective sweep. Part 2 of the proposition, which we prove in Section 6, shows that lineages typically do not coalesce when the favorable B allele dies out. The probability that a selective sweep occurs, and therefore Part 1 of the proposition applies, is $s/(1-(1-s)^{2N})$ (see Durrett (2002) or Schweinsberg and Durrett (2004)).

2.2 A model with recurrent beneficial mutations

To model a population in which beneficial mutations can occur repeatedly, we assume that beneficial mutations at different points on the chromosome occur at times of a Poisson process. The selective advantage that these mutations provide and the rate of recombination between the site of interest and the site of the mutation will be random. When there is a beneficial mutation in the population, the population will evolve as described in the previous subsection. Between these times, the population will follow the standard Moran model.

To be more precise, we will consider the chromosome to be the line segment $[-L, L]$. Our goal will be to describe the genealogy of the site 0. For each N, the beneficial mutations will be governed by a Poisson process K_N on $\mathbb{R} \times [-L, L] \times [0, 1]$. If (t, x, s) is a point in K_N, then at time t, a mutation, which provides a selective advantage of s, will appear at location x on one of the $2N$ chromosomes. The intensity measure of K_N will be $\lambda \times \mu_N$, where λ denotes Lebesgue measure on $\mathbb{R}$ and μ_N is a finite measure on $[-L, L] \times [0, 1]$ which governs the rates of beneficial mutations. The recombination probabilities will be determined by a function $r_N : [-L, L] \to [0, 1]$. We assume that $r_N(0) = 0$ and r_N is nonincreasing on $[-L, 0]$ and nondecreasing on $[0, L]$. Beginning at time t, the population will evolve according to the model described in the previous subsection of a population with a beneficial allele having selective advantage s and recombination probability $r_N(x)$. We let $\tau(t)$ denote the first time that the beneficial mutation that appears at time t either disappears from the population or is present on all $2N$ chromosomes.

Let $\mathcal{T}_N = \{t : (t, x, s) \text{ is a point in } K_N \text{ for some } x \text{ and } s\}$ be the times at which beneficial mutations are proposed. Note, however, that we can not define the evolution of the population as explained above if, for some $t_1, t_2 \in \mathcal{T}_N$, the intervals $[t_1, \tau(t_1)]$ and $[t_2, \tau(t_2)]$ overlap. There has been some work in the biology literature on the question of how a selective sweep is affected by another selective sweep happening at the same time (see, for example, Barton (1995), Gerrish and Lenski (1998), and Kim and Stephen (2003)). However, as we will show, in our model this overlap occurs too infrequently to have any affect on our results, so we avoid the issue of defining the population during periods of overlap by allowing a new beneficial mutation to occur only when there is no other beneficial mutation currently in the population. That is, beneficial mutations will occur at the times in $\mathcal{T}_N' = \{t \in \mathcal{T}_N : \tau(u) < t \text{ for all } u \in \mathcal{T}_N \text{ such that } u < t\}$. Let

$$\mathcal{I}_N = \bigcup_{t \in \mathcal{T}_N'} [t, \tau(t)].$$

A beneficial mutation will be present in the population at time u if and only if $u \in \mathcal{I}_N$. For the intervals in $\mathcal{I}_N$, the evolution of the population was defined in subsection 2.1. For the times in $\mathbb{R} \setminus \mathcal{I}_N$, we will say that the population evolves according to the standard Moran model so that the evolution of the population is well-defined for all of $\mathbb{R}$.

To define the ancestral process $\Psi_N = (\Psi_N(t), t \geq 0)$, we sample n of the $2N$ individuals at random from the population at time zero. We then define $\Psi_N(t)$ to be the partition of

$\{1, \ldots, n\}$ such that i and j are in the same block of $\Psi_N(t)$ if and only if the ith and jth individuals in the sample got their allele at location 0 on the chromosome from the same ancestor at time $-Nt$. Note that we are again speeding up time by a factor of N so that, if there are no beneficial mutations (i.e. if μ_N is the zero measure), the ancestral process $\Psi_N = (\Psi_N(t), t \geq 0)$ is Kingman's coalescent. When we do have beneficial mutations, the ancestral processes will converge as $N \to \infty$, under suitable conditions, to a coalescent with multiple collisions.

2.3 The main convergence theorem and examples

Pitman (1999) introduced coalescents with multiple collisions, in which many blocks of the partition can merge into one. These coalescent processes are in one-to-one correspondence with finite measures Λ on $[0, 1]$, and the coalescent process associated with a particular measure Λ is called the Λ-coalescent. We will consider here only $\mathcal{P}_n$-valued coalescents because they are what we will need to approximate the genealogy of a sample of size n. However, the constructions can be extended, using Kolmogorov's Extension Theorem, to yield coalescent processes that take their values in the set of partitions of $\mathbb{N} = \{1, 2, \ldots\}$.

Suppose $(\Pi_n(t), t \geq 0)$ is the $\mathcal{P}_n$-valued Λ-coalescent. Then $\Pi_n(0)$ is the partition of $\{1, \ldots, n\}$ into singletons. If $\Pi_n(t)$ has b blocks, then every possible transition involves merging k of the blocks into one, where $2 \leq k \leq b$. Denoting the rate of this transition by $\lambda_{b,k}$, we have

$$\lambda_{b,k} = \int_0^1 x^{k-2}(1-x)^{b-k}\, \Lambda(dx). \tag{2.1}$$

If $\Lambda = \delta_0$, where δ_0 denotes a unit mass at zero, then every transition that involves two blocks merging into one happens at rate one, and no other transitions are possible. Thus, the δ_0-coalescent is Kingman's coalescent.

The theorem below states that when we do have beneficial mutations, the ancestral processes converge as $N \to \infty$, under suitable conditions, to a coalescent with multiple collisions. The multiple mergers happen at times of selective sweeps. Note that the convergence is in the sense of finite-dimensional distributions. Convergence in the stronger Skorohod topology does not hold because, during the short time intervals when selective sweeps are taking place, Ψ_N may undergo multiple transitions.

Theorem 2.2. *Let μ be a finite measure on $[-L, L] \times [0, 1]$, and let $r : [-L, L] \to [0, \infty)$ be a bounded continuous function such that $r(0) = 0$ and r is nonincreasing on $[-L, 0]$ and nondecreasing on $[0, L]$. Suppose that, as $N \to \infty$, the measures $N\mu_N$ converge weakly to μ and the functions $(\log 2N)r_N$ converge uniformly to r. Let η be the measure on $(0, 1]$ such that*

$$\eta([y, 1]) = \int_{-L}^{L} \int_0^1 s 1_{\{e^{-r(x)/s} \geq y\}}\, \mu(dx \times ds)$$

for all $y \in (0, 1]$. Let Λ be the measure on $[0, 1]$ defined by $\Lambda = \delta_0 + \Lambda_0$, where $\Lambda_0(dx) = x^2\eta(dx)$. Let $\Pi = (\Pi(t), t \geq 0)$ be the $\mathcal{P}_n$-valued Λ-coalescent. Then, as $N \to \infty$, the finite-dimensional distributions of Ψ_N converge to the finite-dimensional distributions of Π.

Note that in Theorem 2.2, the recombination probability is $O(1/(\log N))$. The function r is assumed to be monotone on $[-L, 0]$ and $[0, L]$ because the greater the distance between 0 and the

site of the mutation, the greater the likelihood of recombination between the two sites. Also, the rate of beneficial mutations is $O(1/N)$, so that the multiple mergers caused by selective sweeps and the ordinary mergers of two lineages at a time are happening on the same time scale. If the rate of selective sweeps were $o(1/N)$, then the multiple mergers would disappear in the limit. If selective sweeps occurred on a faster time scale than $O(1/N)$, then the multiple mergers would dominate for large N and the limiting coalescent would have no δ_0 component. Gillespie (2000) considers this possibility and proposes that it may explain why observed genetic variation does not appear to be as sensitive to population size as Kingman's coalescent model predicts. However, in this paper we focus on the case in which both types of mergers happen on the same time scale.

We now derive the limiting coalescent with multiple collisions in two natural examples.

Example 2.3. Consider the case in which we are concerned only with mutations at a single site, all of which have the same selective advantage. Fix $\alpha > 0$, and let $\mu_N = \alpha N^{-1}\delta_{(z,s)}$ for some $s \in (0,1]$ and $z \in [-L, L]$. This means that beneficial mutations that provide selective advantage s appear on the chromosome at site z at times of a Poisson process. The measures $N\mu_N$ converge to $\mu = \alpha\delta_{(z,s)}$. Assume that the recombination functions r_N are defined such that the sequence $(\log 2N)r_N$ converges uniformly to r, and let $\beta = r(z)$. Then, for all $y \in (0,1]$, we have

$$\eta([y,1]) = \int_{-L}^{L}\int_0^1 u1_{\{e^{-r(x)/u}\geq y\}}\,\mu(dx\times du) = s\alpha 1_{\{e^{-\beta/s}\geq y\}}.$$

Therefore, η consists of a mass $s\alpha$ at $p = e^{-\beta/s}$. It follows from Theorem 2.2 that the limiting coalescent process is the Λ-coalescent, where $\Lambda = \delta_0 + s\alpha p^2\delta_p$. Thus, in addition to the mergers involving just two blocks, we have coalescence events at times of a Poisson process in which we flip p-coins for each lineage and merge the lineages whose coins come up heads.

Example 2.4. It is also natural to consider the case in which mutations occur uniformly along the chromosome. For simplicity, we will assume that the selective advantage s is fixed. Let λ denote Lebesgue measure on $[-L, L]$. Suppose $\mu_N = N^{-1}(\alpha\lambda \times \delta_s)$, so the measures $N\mu_N$ converge to $\mu = \alpha\lambda \times \delta_s$. To model recombination occurring uniformly along the chromosome, we assume that the functions $(\log 2N)r_N$ converge uniformly to the function $r(x) = \beta|x|$, so the probability of recombination is proportional to the distance between the two sites on the chromosome. For all $y \in (0,1]$, we have

$$\eta([y,1]) = \alpha s\int_{-L}^{L} 1_{\{e^{-r(x)/s}\geq y\}}\,dx = \alpha s\int_{-L}^{L} 1_{\{e^{-\beta|x|/s}\geq y\}}\,dx.$$

Since $e^{-\beta|x|/s} \geq y$ if and only if $|x| \leq -(s/\beta)(\log y)$, we have

$$\eta([y,1]) = \min\left\{\frac{-2\alpha s^2\log y}{\beta},\ 2\alpha sL\right\}.$$

Therefore, for $y \geq e^{-\beta L/s}$, we have

$$\frac{d}{dy}\eta([y,1]) = -\frac{2\alpha s^2}{\beta y}.$$

Let $c = 2\alpha s^2/\beta$. It follows that η has a density given by $g_L(y) = c/y$ for $e^{-\beta L/s} \leq y \leq 1$ and $g_L(y) = 0$ otherwise. By Theorem 2.2, the finite-dimensional distributions of the ancestral

processes Ψ_N converge to those of the Λ-coalescent, where $\Lambda = \delta_0 + \Lambda_0$ and Λ_0 has density $h_L(y) = y^2 g_L(y)$. Note that as $L \to \infty$, the density $h_L(y)$ converges to $h(y)$, where $h(y) = cy$ for $y \in [0,1]$ and $h(y) = 0$ otherwise. We can think of this as the limiting coalescent for an infinitely long chromosome.

Example 2.5. Finally, we show that any Λ-coalescent with a unit mass at zero can arise as a limit of ancestral processes in this model. We first show how to obtain coalescents of the form $\Lambda = \delta_0 + \Lambda_0$, where Λ_0 is a finite measure on $[\epsilon, 1]$ and $0 < \epsilon < 1$. Note that in Theorem 2.2, we have $\Lambda_0(dx) = x^2\eta(dx)$, so it suffices to show that μ and r can be chosen to make η an arbitrary finite measure on $[\epsilon, 1]$. Let $G : [\epsilon, 1] \to [0, \infty)$ be any nonincreasing left-continuous function. We will choose μ and r so that $\eta([y,1]) = G(y)$ for $\epsilon \le y \le 1$ and $\eta([0,\epsilon)) = 0$. Let $L = -\frac{1}{2}\log\epsilon$, and let ν be the measure on $[-L, L]$ such that $\nu([-L, 0)) = 0$ and, for $\epsilon \le y \le 1$, $\nu([0, -\frac{1}{2}\log y]) = 2G(y)$. Suppose $r(x) = |x|$ and $\mu = \nu \times \delta_{1/2}$. Then, for $\epsilon \le y \le 1$,

$$\begin{aligned}\eta([y,1]) &= \int_{-L}^{L}\int_0^1 s1_{\{e^{-r(x)/s} \ge y\}}\,\mu(dx \times ds)\\ &= \frac{1}{2}\int_0^L 1_{\{e^{-2x}\ge y\}}\,\nu(dx) = \frac{1}{2}\nu([0, -(\log y)/2]) = G(y),\end{aligned}$$

as claimed. Thus, we can get the Λ-coalescent in the limit if $\Lambda_0((0,\epsilon)) = 0$. We can obtain an arbitrary Λ-coalescent by then taking a limit as $L \to \infty$ (or $\epsilon \downarrow 0$) as in Example 2.4.

3 Approximation by a coalescent with simultaneous multiple collisions

A key ingredient in the proof of Theorem 2.2 is part 1 of Proposition 2.1. Part 1 of Proposition 2.1 says that, up to an error of $O(1/(\log N))$, we can approximate the effect of a selective sweep on the genealogy by flipping a p-coin for each lineage and merging the lineages whose coins come up heads. However, Durrett and Schweinsberg (2004a) observed in simulations that for N between 10,000 and 1,000,000, the approximation in Proposition 2.1 works poorly, largely because it is possible for multiple groups of lineages to coalesce at the time of a selective sweep. By taking this into account, they were able to give a more complicated approximation that works much better in simulations and has an error of only $O(1/(\log N)^2)$.

Before stating this result, we review Kingman's (1978) paintbox construction of exchangeable random partitions of $\{1, \dots, n\}$. Let

$$\Delta = \Big\{(x_1, x_2, \dots) : x_1 \ge x_2 \ge \cdots \ge 0, \sum_{i=1}^{\infty} x_i \le 1\Big\},$$

and let G be a probability measure on Δ. We define a G-partition Π of $\{1, \dots, n\}$ as follows. Let $Y = (Y_1, Y_2, \dots)$ be a Δ-valued random variable with distribution G. Define a sequence $(Z_i)_{i=1}^n$ to be conditionally i.i.d. given Y such that $P(Z_i = j|Y) = Y_j$ for all positive integers j and $P(Z_i = 0|Y) = 1 - \sum_{j=1}^{\infty} Y_j$. Then define Π to be the partition such that distinct integers i and j are in the same block if and only if $Z_i = Z_j \ge 1$. We denote the distribution of a G-partition of $\{1, \dots, n\}$ by $Q_{G,n}$. Note that if G is a unit mass at $(p, 0, 0, \dots)$, then $Q_{G,n} = Q_{p,n}$.

Next, we define a family of distributions $R(\theta, M)$ on Δ by using a stick-breaking construction. Let $\theta \in [0,1]$, and let M be a positive integer. Let $(W_k)_{k=2}^M$ be independent random variables such that W_k has a Beta$(1, k-1)$ distribution. Let $(\zeta_k)_{k=2}^M$ be a sequence of independent random variables such that $P(\zeta_k = 1) = \theta$ and $P(\zeta_k = 0) = 1 - \theta$ for all k. For $k = 2, 3, \ldots, M$, let $V_k = \zeta_k W_k$. To perform the stick breaking, we first break off a fraction W_M of the unit interval, then break off a fraction W_{M-1} of what is left over, and so on until we get down to W_2. For $k = 2, \ldots, M$, the length of the kth fragment is $\tilde{Y}_k = V_k \prod_{j=k+1}^M (1-V_j)$, and the length of the first fragment is $\tilde{Y}_1 = \prod_{j=2}^M (1 - V_j)$. Note that $\sum_{k=1}^M \tilde{Y}_k = 1$. Let $Y = (Y_1, Y_2, \ldots, Y_M, 0, 0, \ldots) \in \Delta$ be the sequence obtained by ranking the interval lengths $\tilde{Y}_1, \ldots, \tilde{Y}_M$ in decreasing order and then appending an infinite sequence of zeros. Finally, let $R(\theta, M)$ be the distribution of Y.

These distributions $R(\theta, M)$ were studied in Durrett and Schweinsberg (2004b), who used them to approximate the distribution of family sizes in a Yule process with infinitely many types. They arise in the proposition below because, after a beneficial mutation, the number of lineages with the B allele that do not eventually die out can be approximated by a Yule process. The result below is Theorem 1.2 of Schweinsberg and Durrett (2004).

Proposition 3.1. *Fix $n \in \mathbb{N}$, and fix $s \in (0,1)$. Assume there is a constant C' such that $r \leq C'/(\log N)$ for all N. Let $\alpha = r \log(2N)/s$, and let $p = e^{-\alpha}$. Then there exists a positive constant C, depending continuously on s and α but not depending on N, such that*

$$|P(\Theta = \pi | X(\tau) = 2N) - Q_{R(r/s, \lfloor 2Ns \rfloor), n}(\pi)| \leq C/(\log N)^2$$

for all $\pi \in \mathcal{P}_n$, where $\lfloor m \rfloor$ denotes the greatest integer less than or equal to m.

Because the improved approximation allows many groups of lineages to coalesce at the time of a selective sweep, this result suggests that, for finite N, a coalescent with simultaneous multiple collisions should provide a better approximation of the ancestral process than a coalescent with multiple collisions. Coalescents with simultaneous multiple collisions, which were studied by Möhle and Sagitov (2001), Schweinsberg (2000), and Bertoin and Le Gall (2003), have the property that many blocks can merge at once into a single block, and many such mergers can occur simultaneously. Coalescents with simultaneous multiple collisions are in one-to-one correspondence with finite measures Ξ on Δ.

Suppose π is a partition of $\{1, \ldots, n\}$ whose blocks are $B_1, \ldots, B_m$, and suppose π' is a partition of $\{1, \ldots, n'\}$ with $n' \geq m$ whose blocks are $B_1', \ldots, B_k'$. Following Bertoin and Le Gall (2003), define the coagulation of π by π' to be the partition whose blocks are given by $\bigcup_{j \in B_i'} B_j$ for $i = 1, \ldots, k$. Suppose $(\Pi_n(t), t \geq 0)$ is the $\mathcal{P}_n$-valued Ξ-coalescent. If there are b blocks at time $t-$ and a merger occurs at time t, then there exists a unique partition $\pi \in \mathcal{P}_b$ such that $\Pi_n(t)$ is the coagulation of $\Pi_n(t-)$ by π. If π has $r + s$ blocks, s of which are singletons and the other r of which have sizes $k_1, \ldots, k_r \geq 2$, where $b = k_1 + \cdots + k_r + s$, then the rate of this transition is

$$\lambda_{b;k_1,\ldots,k_r;s} = \int_\Delta Q_{\delta_x, b}(\pi) \left(\sum_{j=1}^\infty x_j^2 \right)^{-1} \Xi_0(dx) + a 1_{\{r=1, k_1=2\}}, \tag{3.1}$$

where δ_x denotes a unit mass at $x = (x_1, x_2, \ldots) \in \Delta$ and Ξ has been written as $a\delta_{(0,0,\ldots)} + \Xi_0$ with $\Xi_0(\{(0,0,\ldots)\}) = 0$. Coalescents with multiple collisions are a special case in which Ξ is concentrated on points in which only the first coordinate is nonzero.

Coalescents with multiple and simultaneous multiple collisions can be constructed from Poisson point processes (see Pitman (1999) and Schweinsberg (2000)). Consider a Poisson process on $(0,\infty)\times\mathcal{P}_n$ whose intensity measure is the product of Lebesgue measure on $(0,\infty)$ and a measure L on $\mathcal{P}_n$ defined as follows. Let $S\subset\mathcal{P}_n$ be the set of all partitions consisting of one block of size 2 and $n-2$ singletons. If $\pi\in\mathcal{P}_n$, let $L(\pi)=0$ if π is the partition consisting of n singletons. Otherwise, let

$$L(\pi)=\int_\Delta Q_{\delta_x,n}(\pi)\biggl(\sum_{j=1}^\infty x_j^2\biggr)^{-1}\,\Xi_0(dx)+a1_{\{\pi\in S\}}. \tag{3.2}$$

Since L is a finite measure, it is easy to define $\Pi_n=(\Pi_n(t),t\ge 0)$ such that $\Pi_n(0)$ is the partition consisting of n singletons and, at the times of points (t,π) of the Poisson point process, the partition $\Pi_n(t)$ is the coagulation of $\Pi_n(t-)$ by π, and these are the only jump times of Π_n. This coalescent process is the $\mathcal{P}_n$-valued Ξ-coalescent. The construction of the Λ-coalescent is the same, except that if π has at least one block that is not a singleton, we define

$$L(\pi)=\int_0^1 Q_{p,n}(\pi)p^{-2}\,\Lambda_0(dp)+a1_{\{\pi\in S\}}, \tag{3.3}$$

where $\Lambda=\delta_0+\Lambda_0$ and $\Lambda_0(\{0\})=0$.

Under some additional assumptions, most significantly restricting the selective advantage resulting from each beneficial mutation to be at least $\epsilon>0$, we are able to obtain bounds on the difference between the finite-dimensional distributions of Ψ_N and the finite-dimensional distributions of the approximating coalescent process. Proposition 3.2 below shows that indeed the coalescent with simultaneous multiple collisions gives a more accurate approximation.

Proposition 3.2. *Let μ be a finite measure on $[-L,L]\times[\epsilon,1]$, where $\epsilon>0$, and let $r:[-L,L]\to[0,1]$ be a function such that $r(0)=0$ and r is nonincreasing on $[-L,0]$ and nondecreasing on $[0,L]$. Suppose that, for all N, we have $\mu_N=N^{-1}\mu$. Also, assume that $r_N(x)=r(x)/\log(2N)$ for all N and x. Fix times $0<u_1<\cdots<u_m$, and let $\pi_1,\ldots,\pi_m\in\mathcal{P}_n$.*

1. *Define η and Λ as in Theorem 2.2. Let $\Pi=(\Pi(t),t\ge 0)$ be the $\mathcal{P}_n$-valued Λ-coalescent. Then there exists a constant C such that*

$$|P(\Psi_N(u_i)=\pi_i \text{ for } i=1,\ldots,m)-P(\Pi(u_i)=\pi_i \text{ for } i=1,\ldots,m)|\le\frac{C}{\log N}.$$

2. *Let G_N be the measure on Δ such that for all measurable subsets $A\subset\Delta$, we have*

$$G_N(A)=\int_{-L}^L\int_0^1 sR(r_N(x)/s,\lfloor 2Ns\rfloor)(A)\,\mu(dx\times ds).$$

Let Ξ_N be the measure on Δ given by $\Xi_N=\delta_{(0,0,\ldots)}+\Xi_{N,0}$, where $\Xi_{N,0}$ is defined by $\Xi_{N,0}(dx)=(\sum_{j=1}^\infty x_j^2)G_N(dx)$. Let $\Upsilon_N=(\Upsilon_N(t),t\ge 0)$ be the $\mathcal{P}_n$-valued Ξ_N-coalescent. Then there exists a constant C such that

$$|P(\Psi_N(u_i)=\pi_i \text{ for } i=1,\ldots,m)-P(\Upsilon_N(u_i)=\pi_i \text{ for } i=1,\ldots,m)|\le\frac{C}{(\log N)^2}.$$

4 Segregating sites and pairwise differences

One motivation for modeling a population that experiences recurrent selective sweeps by coalescents with multiple or simultaneous multiple collisions is that these coalescent models can provide insight into tests used to detect selective sweeps. In view of part 2 of Proposition 3.2 and the simulation results in Durrett and Schweinsberg (2004a), there should be little loss of accuracy in studying the behavior of these tests under the assumption that the genealogy of a sample follows a coalescent with simultaneous multiple collisions. One commonly used test is based on Tajima's D-statistic (see Tajima (1989)). Given a sample of n strands of DNA from the same region on a chromosome, let Δ_{ij} be the number of sites at which the ith and jth segments differ, and let $\Delta_n = \binom{n}{2}^{-1} \sum_{i \neq j} \Delta_{ij}$ be the average number of pairwise differences over the $\binom{n}{2}$ possible pairs. Let S_n be the number of segregating sites in the sample, that is, the number of sites at which at least one pair of segments differs. Tajima's D-statistic compares the statistics Δ_n and S_n.

Suppose the ancestral history of a sample of N individuals is given by a coalescent with multiple or simultaneous multiple collisions. Let λ_b be the total rate of all mergers when the coalescent has b blocks. Assume that, on the time scale of the coalescent process, mutations happen at rate $\theta/2$. Any mutation on the ith or jth lineage before these lineages coalesce will cause the ith and jth segments to differ at some site. Since the expected time for these lineages to coalesce is λ_2^{-1}, we have $E[\Delta_{ij}] = \theta \lambda_2^{-1}$. Therefore

$$E[\Delta_n] = \theta \lambda_2^{-1}. \tag{4.1}$$

Note that $\lambda_2 = \Lambda([0,1])$ for coalescents with multiple collisions and $\lambda_2 = \Xi(\Delta)$ for coalescents with simultaneous multiple collisions.

To calculate the expected number of segregating sites, we note that any mutation in the ancestral tree before all n lineages have coalesced into one adds to the number of segregating sites. If, at some time, the coalescent has exactly b blocks, the expected time that the coalescent has b blocks is λ_b^{-1}. Let $G_n(b)$ be the probability that the coalescent, starting with n blocks, will have exactly b blocks at some time. Then

$$E[S_n] = \frac{\theta}{2} \sum_{b=2}^{n} b \lambda_b^{-1} G_n(b). \tag{4.2}$$

Although we do not have a closed-form expression for $G_n(b)$, these quantities can be calculated recursively because (2.1) and (3.1) allow us to express $G_n(b)$ in terms of $G_k(b)$ for $k < n$. As a result, it would not be difficult to evaluate the expression in (4.2) numerically.

Suppose the ancestral process is given by Kingman's coalescent, which would be the case if there were no selective sweeps. Then $\lambda_b = \binom{b}{2}$ for all $b \geq 2$. Also, the number of blocks never decreases by more than one at a time, so $G_n(b) = 1$ whenever $2 \leq b \leq n$. It follows that $E[\Delta_n] = \theta$ and

$$E[S_n] = \frac{\theta}{2} \sum_{b=2}^{n} b \binom{b}{2}^{-1} = \theta \sum_{b=2}^{n} \frac{1}{b-1} = \theta h_{n-1}, \tag{4.3}$$

where $h_{n-1} = \sum_{i=1}^{n-1} (1/i)$. Thus, $E[\Delta_n - S_n/h_{n-1}] = 0$. This observation is the basis for Tajima's D-statistic, which is given by

$$D = \frac{\Delta_n - S_n/h_{n-1}}{\sqrt{a_n S_n + b_n S_n (S_n - 1)}}, \tag{4.4}$$

where a_n and b_n are somewhat complicated constants that are chosen to make the variance of D approximately one when the ancestral tree is given by Kingman's coalescent. See section 4.1 of Durrett (2002) for details.

After a selective sweep, the new mutants will tend to have low frequency. As a result, a recent selective sweep should decrease Δ_n more than S_n, causing the numerator of Tajima's D-statistic to be negative. Braverman et. al. (1995) found in simulations that Tajima's D-statistic indeed tends to be negative after a selective sweep. Simonsen, Churchill, and Aquadro (1995) studied this question further and argued that unless the selective sweep was recent, Tajima's D-statistic had relatively little power to detect selective sweeps. See also Przeworski (2002), who discusses the power of Tajima's D-statistic to detect selective sweeps. Our coalescent approximation allows us to obtain the following result regarding the expected number of segregating sites when the population experiences recurrent selective sweeps.

Proposition 4.1. *Consider a Λ-coalescent in which $\Lambda = \delta_0 + \Lambda_0$, where $\Lambda_0(\{0\}) = 0$, or a Ξ-coalescent in which $\Xi = \delta_{(0,0,\dots)} + \Xi_0$ and $\Xi_0(\{(0,0,\dots)\}) = 0$. Let $\alpha_b = \lambda_b - \binom{b}{2}$. Suppose*

$$\sum_{b=2}^{\infty} \frac{\alpha_b \log b}{b^2} < \infty. \tag{4.5}$$

Then, there exists a constant $\rho \geq 0$ such that

$$\lim_{n\to\infty} E[S_n] - \theta h_{n-1} = -\rho. \tag{4.6}$$

Furthermore, defining $G_\infty(b) = \lim_{n\to\infty} G_n(b)$, we have

$$\rho = \frac{\theta}{2} \sum_{b=2}^{\infty} b\left(\binom{b}{2}^{-1} - \lambda_b^{-1} \right) + \frac{\theta}{2} \sum_{b=2}^{\infty} b\lambda_b^{-1}(1 - G_\infty(b)). \tag{4.7}$$

The condition (4.5) prevents Λ_0 or Ξ_0 from having too much mass near zero. Note that (4.1) implies that $E[\Delta_n]$ decreases by a constant as a result of the beneficial mutations, while Proposition 4.1 implies that when (4.5) holds, $E[S_n/h_{n-1}]$ decreases by approximately ρ/h_{n-1}, which is $O(1/(\log n))$. Therefore, Proposition 4.1 shows that for sufficiently large samples we do expect Tajima's D-statistic to be negative when the population is affected by recurrent selective sweeps. Before proving this proposition, we consider some examples.

Example 4.2. Suppose, as in Example 2.3, we have a Λ-coalescent in which $\Lambda = \delta_0 + s\alpha p^{-2}\delta_p$. Since p-mergers occur at rate $s\alpha$, we have $\lambda_b \leq \binom{b}{2} + s\alpha$ and thus $\alpha_b \leq s\alpha$ for all b. Condition (4.5) follows immediately.

Suppose instead we have the Λ-coalescent of Example 2.4, where $\Lambda = \delta_0 + \Lambda_0$ and $\Lambda_0(dx) = cx\,dx$. Note that α_b is the same as the total merger rate of the Λ_0-coalescent when there are b blocks. Using the fact that if $Z \sim \text{Binomial}(b, x)$ then $P(Z \geq 2) = 1 - (1-x)^b - bx(1-x)^{b-1}$,

we have

$$\begin{aligned}
\alpha_b &= \int_0^1 (1-(1-x)^b - bx(1-x)^{b-1})x^{-2}\,\Lambda_0(dx) \\
&= c\int_0^1 (1-(1-x)^b - bx(1-x)^{b-1})x^{-1}\,dx \le c\int_0^1 (1-(1-x)^b)x^{-1}\,dx \\
&= c\int_0^{1/b} (1-(1-x)^b)x^{-1}\,dx + c\int_{1/b}^1 (1-(1-x)^b)x^{-1}\,dx \\
&\le c\int_0^{1/b} b\,dx + c\int_{1/b}^1 x^{-1}\,dx = c(1+\log b),
\end{aligned} \tag{4.8}$$

which implies (4.5).

Example 4.3. Although (4.5) holds in the natural cases given in Examples 2.3 and 2.4, we show here that it does not hold for all coalescents. Suppose $\Lambda = \delta_0 + \Lambda_0$, where Λ_0 is the uniform distribution on $(0,1)$. Note that there exists a constant $C > 0$ such that if $Z \sim \text{Binomial}(b,x)$ with $x \ge 1/b$ and $b \ge 2$, then $P(Z \ge 2) \ge C$. Therefore,

$$\begin{aligned}
\alpha_b &= \int_0^1 (1-(1-x)^b - bx(1-x)^{b-1})x^{-2}\,dx \\
&\ge \int_{1/b}^1 (1-(1-x)^b - bx(1-x)^{b-1})x^{-2}\,dx \\
&\ge C\int_{1/b}^1 x^{-2}\,dx = C(b-1),
\end{aligned}$$

so (4.5) does not hold in this case.

Proof of Proposition 4.1. When the coalescent has $n+1$ blocks, the probability that the next coalescence event will take the coalescent down to fewer than n blocks is at most $[\lambda_{n+1} - \binom{n+1}{2}]/\lambda_{n+1}$. Therefore, if $2 \le b \le n$, then

$$|G_{n+1}(b) - G_n(b)| \le \frac{\lambda_{n+1} - \binom{n+1}{2}}{\lambda_{n+1}} = \frac{\alpha_{n+1}}{\lambda_{n+1}} \le \frac{2\alpha_{n+1}}{n(n+1)}. \tag{4.9}$$

Therefore, when (4.5) holds, the sequence $(G_n(b))_{n=b}^\infty$ is Cauchy and thus has a limit $G_\infty(b)$.

It follows from (4.2) and (4.3) that

$$\begin{aligned}
E[S_n] - \theta h_{n-1} &= \frac{\theta}{2}\sum_{b=2}^n b\lambda_b^{-1}G_n(b) - \frac{\theta}{2}\sum_{b=2}^n b\binom{b}{2}^{-1} \\
&= \frac{\theta}{2}\sum_{b=2}^n b\left(\lambda_b^{-1} - \binom{b}{2}^{-1}\right) - \frac{\theta}{2}\sum_{b=2}^n b\lambda_b^{-1}(1-G_\infty(b)) \\
&\qquad + \frac{\theta}{2}\sum_{b=2}^n b\lambda_b^{-1}(G_n(b) - G_\infty(b)).
\end{aligned} \tag{4.10}$$

To prove Proposition 4.1, we need to take the limit as $n \to \infty$ of the three terms on the right-hand side of (4.10).

For the first term, we note that

$$\binom{b}{2}^{-1} - \lambda_b^{-1} = \frac{\lambda_b - \binom{b}{2}}{\binom{b}{2}\lambda_b} \le \alpha_b \binom{b}{2}^{-2} = \frac{4\alpha_b}{b^2(b-1)^2}.$$

Therefore, when (4.5) holds, we have a summable series and

$$\lim_{n\to\infty} \frac{\theta}{2} \sum_{b=2}^{n} b\left(\lambda_b^{-1} - \binom{b}{2}^{-1}\right) = -\frac{\theta}{2} \sum_{b=2}^{\infty} b\left(\binom{b}{2}^{-1} - \lambda_b^{-1}\right). \tag{4.11}$$

For the second term, note that (4.9) and the fact that $G_b(b) = 1$ imply

$$\begin{aligned}\sum_{b=2}^{\infty} b\lambda_b^{-1}(1 - G_\infty(b)) &\le \sum_{b=2}^{\infty} \frac{2}{b-1}\left(\sum_{m=b}^{\infty} \frac{2\alpha_{m+1}}{m(m+1)}\right) \\ &= \sum_{m=2}^{\infty} \frac{2\alpha_{m+1}}{m(m+1)} \sum_{b=2}^{m} \frac{2}{b-1} \le \sum_{m=2}^{\infty} \frac{4\alpha_{m+1}(1+\log(m-1))}{m(m+1)},\end{aligned}$$

which is finite by (4.5). Therefore,

$$\lim_{n\to\infty} \frac{\theta}{2} \sum_{b=2}^{n} b\lambda_b^{-1}(1 - G_\infty(b)) = \frac{\theta}{2} \sum_{b=2}^{\infty} b\lambda_b^{-1}(1 - G_\infty(b)). \tag{4.12}$$

Finally, for the third term,

$$\begin{aligned}\limsup_{n\to\infty} \sum_{b=2}^{n} b\lambda_b^{-1}|G_n(b) - G_\infty(b)| &\le \limsup_{n\to\infty} \sum_{b=2}^{n} \frac{2}{b-1}\left(\sum_{m=n}^{\infty} \frac{2\alpha_{m+1}}{m(m+1)}\right) \\ &\le \limsup_{n\to\infty} \frac{1}{\log n} \sum_{b=2}^{n} \frac{2}{b-1}\left(\sum_{m=n}^{\infty} \frac{2\alpha_{m+1}\log m}{m(m+1)}\right) \\ &\le \limsup_{n\to\infty} \frac{2(1+\log(n-1))}{\log n} \sum_{m=n}^{\infty} \frac{2\alpha_{m+1}\log m}{m(m+1)} = 0 \end{aligned} \tag{4.13}$$

by (4.5). The proposition follows from (4.10), (4.11), (4.12), and (4.13). □

5 The number of singletons

Fu and Li (1993) proposed another test to detect departures from Kingman's coalescent. They considered the ancestral tree in which the leaves are the n individuals in the sample. They defined the branches connecting a leaf to an internal node to be external branches and the other branches to be internal branches. Let η_e denote the number of mutations on external branches, and let η_i be the number of mutations on internal branches. Every mutation produces a segregating site, so $\eta_e + \eta_i = S_n$. If a mutation occurs on an external branch, the mutant gene appears on just one of the n individuals in the sample, while if a mutation occurs on an internal branch, the mutant gene appears on between 2 and $n-1$ of the individuals in the sample. Therefore, to determine η_e, we simply count the number of mutations that appear on just one of the sampled chromosomes.

Note that unless an outgroup is available, it will not be possible to distinguish between a mutation that appears on one of the sampled chromosomes and a mutation that appears on $n-1$ of the sampled chromosomes. Fu and Li (1993) proposed a modification of their test for when there is no outgroup, but for the analysis in this section, we assume that we have an outgroup that enables us to make this distinction.

Let J_n be the sum of the lengths of the external branches. In terms of the associated coalescent process, J_n is the sum, over i between 1 and n, of the amount of time that the integer i is in a singleton block. Let I_n be the sum of the lengths of the internal branches. Assuming, as before, that mutations occur at rate $\theta/2$ on the time scale of the coalescent process, we have $E[\eta_e|J_n] = (\theta/2)J_n$ and $E[\eta_i|I_n] = (\theta/2)I_n$.

Fu and Li's D-statistic is based on comparing η_i with $(h_{n-1}-1)\eta_e$. Note that $\eta_i-(h_{n-1}-1)\eta_e = S_n - h_{n-1}\eta_e$. To see that this has mean zero when the ancestral tree is given by Kingman's coalescent, we follow the explanation on p. 163 of Durrett (2002). In the case of Kingman's coalescent, (4.3) gives $E[S_n] = \theta h_{n-1}$. Therefore, $E[S_n - h_{n-1}\eta_e] = \theta h_{n-1} - \theta h_{n-1}E[J_n]/2$, so it remains to show that $E[J_n] = 2$. Let K_n be the amount of time that the integer 1 is in a singleton block of the partition, so $E[J_n] = nE[K_n]$. Let T_n be the amount of time before the first coalescence event, and note that $E[T_n] = 2/[n(n-1)]$. The probability that 1 coalesces with another integer at time T_n is $2/n$, and this event is independent of T_n. If 1 does not coalesce at this time, then the expected additional time that 1 is a singleton is $E[K_{n-1}]$. Therefore, we get the recursion

$$E[K_n] = \frac{2}{n}E[T_n] + \frac{n-2}{n}E[T_n + K_{n-1}] = \frac{2}{n(n-1)} + \frac{n-2}{n}E[K_{n-1}].$$

Note that $E[K_2] = 1$, and then it is easy to show by induction that $E[K_n] = 2/n$ for all n, and so $E[J_n] = 2$ for all n, as claimed.

We can write Fu and Li's D-statistic as

$$D = \frac{S_n - h_{n-1}\eta_e}{\sqrt{c_n S_n + d_n S_n^2}}, \tag{5.1}$$

where, as in (4.4), c_n and d_n are constants chosen to make the variance of the statistic approximately one when the genealogy is given by Kingman's coalescent. Details of the variance computation are given in section 4.2 in Durrett (2002), where an error of Fu and Li (1993) is corrected.

When multiple mergers cause many lineages to coalesce at once, one expects I_n to be reduced more than J_n because there is still an external branch associated with each leaf, but there are fewer internal branches because of multiple mergers. This would cause Fu and Li's D-statistic to be negative. The next proposition shows that this is indeed the case.

Proposition 5.1. *Let $(\Pi_n(t), t \geq 0)$ be a $\mathcal{P}_n$-valued Λ-coalescent in which $\Lambda = \delta_0 + \Lambda_0$, where $\Lambda_0(\{0\}) = 0$, or a $\mathcal{P}_n$-valued Ξ-coalescent in which $\Xi = \delta_{(0,0,\dots)} + \Xi_0$ and $\Xi_0(\{(0,0,\dots)\}) = 0$. Let $\alpha_b = \lambda_b - \binom{b}{2}$, and suppose (4.5) holds. Then*

$$\lim_{n\to\infty} E[S_n - h_{n-1}\eta_e] = -\rho, \tag{5.2}$$

where ρ is the constant defined in (4.7).

The key to the proof of this proposition is the following lemma.

Lemma 5.2. *Under the assumptions of Proposition 5.1, there is a positive constant C such that*

$$0 \le E[2 - J_n] \le \frac{C}{n} \sum_{b=2}^{n} \frac{\alpha_b}{b} \tag{5.3}$$

for all $n \ge 2$.

The first inequality in (5.3), which does not require condition (4.5), shows that the expected sum of the lengths of the external branches is never greater than 2, which means that it is largest for Kingman's coalescent. The second inequality gives a rather sharp bound on the difference. Recall that in Example 2.3, we have $\alpha_b \le s\alpha$, so $E[2 - J_n] \le C'(\log n)/n$ for some other constant C'. In Example 2.4, (4.8) gives $\alpha_b \le c(1 + \log b) \le c(1 + \log n)$, which implies $E[2 - J_n] \le C''(\log n)^2/n$ for some constant C''. Thus, in these examples, the lengths of the external branches are affected very little by multiple mergers when the sample size is large. The reason is that, in large samples, a lot of coalescence occurs very quickly, so most ancestral lines have merged with at least one other ancestral line before the first multiple merger takes place.

Proof of Lemma 5.2. We start by proving the first inequality in (5.3) by induction. As before, let K_n be the amount of time that the integer 1 is in a singleton block. We need to show that $E[K_n] \le 2/n$ for all $n \ge 2$. First, note that $E[K_2] = \lambda_2^{-1} \le 1$. Now, suppose for some $n \ge 3$, we have $E[K_j] \le 2/j$ for $j = 2, \ldots, n-1$, and consider $E[K_n]$. Let T_n be the time of the first merger when the coalescent starts with n blocks, and let $B \ge 2$ be the number of blocks involved in the merger at time T_n. Note that B is independent of T_n. Conditional on B, the probability that 1 merges with at least one other block at time T_n is B/n. If this does not happen, then at least $n - B + 1$ blocks remain after the merger, so by the induction hypothesis, the expected time after T_n that $\{1\}$ will remain a singleton is at most $2/(n - B + 1)$. Therefore,

$$E[K_n | T_n, B] \le \left(\frac{B}{n}\right) T_n + \left(\frac{n-B}{n}\right)\left(T_n + \frac{2}{n-B+1}\right) = T_n + \frac{2(n-B)}{n(n-B+1)}.$$

Since $2 \le B \le n$, we have $(n-B)/(n-B+1) \le (n-2)/(n-1)$. Also, $E[T_n] = \lambda_n^{-1} \le 2/[n(n-1)]$, so

$$E[K_n] \le \frac{2}{n(n-1)} + \frac{2(n-2)}{n(n-1)} = \frac{2}{n},$$

which proves the first inequality.

The proof of the second inequality requires a coupling argument. Let $(\Pi_n(t), t \ge 0)$ be the coalescent process defined in the statement of Proposition 5.1, and let $(\Upsilon_n(t), t \ge 0)$ be Kingman's coalescent, started from the partition of $1, \ldots, n$ into singletons. We may assume that the coalescent processes Π_n and Υ_n are constructed from Poisson processes N_1 and N_2 respectively on $(0, \infty) \times \mathcal{P}_n$, as described in section 3. That is, whenever (t, π) is a point of N_1, the partition $\Pi_n(t)$ is the coagulation of $\Pi_n(t-)$ by π, and whenever (t, π) is a point of N_2, the partition $\Upsilon_n(t)$ is the coagulation of $\Upsilon_n(t-)$ by π. Furthermore, these are the only jump times of Π_n and Υ_n. Let L_1 and L_2 be the intensity measures of the second coordinate for the Poisson processes N_1 and N_2 respectively. Then, for $\pi \in \mathcal{P}_n$, we have $L_2(\pi) = 1$ if π consists of one block of size 2 and $n - 2$ singletons, and $L_2(\pi) = 0$ otherwise. Also, $L_1(\pi) \ge L_2(\pi)$ for

all $\pi \in \mathcal{P}_n$. Therefore, we may assume that the Poisson processes N_1 and N_2 are coupled such that if (t,π) is a point of N_2 then (t,π) is a point of N_1. The points (t,π) in both N_1 and N_2 correspond to mergers in which two blocks coalesce at a time, while the points (t,π) in N_1 but not N_2 correspond to multiple mergers caused by selective sweeps.

To compare the two processes, note that $K_n = \inf\{t : \{1\} \text{ is not a singleton in } \Pi_n(t)\}$, and let $K_n' = \inf\{t : \{1\} \text{ is not a singleton in } \Upsilon_n(t)\}$. We have $E[J_n] = nE[K_n]$. By our previous results for Kingman's coalescent, we have $E[K_n'] = 2/n$, and so $E[2 - J_n] = nE[K_n' - K_n]$. Let $\tau = \inf\{t : \Pi_n(t) \neq \Upsilon_n(t)\}$, where we say $\tau = \infty$ if $\Pi_n(t) = \Upsilon_n(t)$ for all t. For $\pi \in \mathcal{P}_n$, denote by $|\pi|$ the number of blocks in π. Since $\Pi_n(t) = \Upsilon_n(t)$ for all $t \leq \tau$, we have

$$E[2 - J_n] = nE[K_n' - K_n] \leq nE[(K_n' - \tau)1_{\{\tau < K_n'\}}] = n\sum_{b=2}^{n} E[(K_n' - \tau)1_{\{\tau < K_n'\}}1_{\{|\Upsilon_n(\tau)|=b\}}].$$

For $b = 1, 2, \ldots, n$, define $T_b = \inf\{t : |\Upsilon_n(t)| = b\}$. If $\tau < K_n'$ and $|\Upsilon_n(\tau)| = b$, then $K_n' > T_b$. Therefore,

$$E[2 - J_n] \leq n\sum_{b=2}^{n} E[K_n' - \tau | \{\tau < K_n'\} \cap \{|\Upsilon_n(\tau)| = b\}] P(\{K_n' > T_b\} \cap \{|\Upsilon_n(\tau)| = b\}). \quad (5.4)$$

If $\tau < K_n'$ and $|\Upsilon_n(\tau)| = b$, then $\{1\}$ is one of b blocks of $\Upsilon_n(\tau)$, and by our previous results on Kingman's coalescent, the expected time before it merges with another block is $2/b$. Thus, we have

$$E[K_n' - \tau | \{\tau < K_n'\} \cap \{|\Upsilon_n(\tau)| = b\}] = \frac{2}{b}. \quad (5.5)$$

Note that $K_n' > T_b$ whenever $\{1\}$ remains a singleton at the time that Kingman's coalescent is down to b blocks. Whenever the coalescent goes from j blocks to $j-1$, the probability that the integer 1 is involved in the merger is $2/j$, so

$$P(K_n' > T_b) = \prod_{j=b+1}^{n} \left(1 - \frac{2}{j}\right) \leq \exp\left(-\sum_{j=b+1}^{n} \frac{2}{j}\right) \leq \exp\left(1 - 2\int_b^n \frac{1}{x}\, dx\right) = e\left(\frac{b}{n}\right)^2. \quad (5.6)$$

If $|\Upsilon_n(\tau)| = b$, then both Π_n and Υ_n have the same b blocks at time T_b, but at time τ the process Π_n has a transition but Υ_n does not. Since the total merger rate for Π_n after time T_b is $\lambda_b = \alpha_b + \binom{b}{2}$ and the total merger rate for Υ_n after time T_b is $\binom{b}{2}$, we have

$$P(|\Upsilon_n(\tau)| = b | K_n' > T_b) \leq \frac{\alpha_b}{\lambda_b} \leq \frac{2\alpha_b}{b(b-1)}. \quad (5.7)$$

Combining (5.4)-(5.7), we get

$$E[2 - J_n] \leq n\sum_{b=2}^{n} \frac{4e\alpha_b b^2}{b^2(b-1)n^2} \leq \frac{C}{n}\sum_{b=2}^{n} \frac{\alpha_b}{b},$$

which is the second inequality in (5.3). □

Proof of Proposition 5.1. We have

$$E[S_n - h_{n-1}\eta_e] = (E[S_n] - \theta h_{n-1}) + h_{n-1}(\theta - E[\eta_e]) = (E[S_n] - \theta h_{n-1}) + \frac{h_{n-1}\theta}{2}E[2 - J_n]. \quad (5.8)$$

By Proposition 4.1, $\lim_{n\to\infty}(E[S_n] - \theta h_{n-1}) = -\rho$. It thus remains only to show that the second term on the right-hand side of (5.8) goes to zero as $n \to \infty$. Let $\epsilon > 0$. By (4.5), there exists a positive integer N such that

$$\sum_{b=N}^{\infty} \frac{\alpha_b(1 + \log b)}{b^2} < \epsilon.$$

Therefore, by Lemma 5.2,

$$\begin{aligned}\limsup_{n\to\infty} \frac{h_{n-1}\theta}{2}E[2 - J_n] &\leq \limsup_{n\to\infty} \frac{Ch_{n-1}\theta}{2n}\sum_{b=2}^{n}\frac{\alpha_b}{b} = \limsup_{n\to\infty} \frac{Ch_{n-1}\theta}{2n}\left(\sum_{b=2}^{N}\frac{\alpha_b}{b} + \sum_{b=N}^{n}\frac{\alpha_b}{b}\right) \\ &\leq 0 + \frac{C\theta}{2}\limsup_{n\to\infty}\sum_{b=N}^{n}\frac{\alpha_b h_{n-1}}{bn} \leq \frac{C\theta}{2}\limsup_{n\to\infty}\sum_{b=N}^{n}\frac{\alpha_b(1+\log b)}{b^2} \leq \frac{C\theta\epsilon}{2}.\end{aligned}$$

Since this is true for all $\epsilon > 0$, and since $E[2 - J_n] \geq 0$ for all n by Lemma 5.2, we have

$$\lim_{n\to\infty} \frac{h_{n-1}\theta}{2}E[2 - J_n] = 0,$$

which completes the proof of the proposition. □

We conclude this section with some comments about the power of Tajima's D-statistic and Fu and Li's D-statistic to detect selective sweeps. The numerators of these two statistics, which are $\Delta_n - S_n/h_{n-1}$ and $S_n - h_{n-1}\eta_e$, each have mean zero when the ancestral process is Kingman's coalescent. The expected values of these two numerators both converge to a negative constant as the sample size goes to infinity when multiple mergers can occur. These statistics are used to test for departures from Kingman's coalescent. If the goal is to test for multiple mergers caused by selective sweeps, one would reject the null hypothesis of no selective sweeps if the value of the statistic is too small (i.e. more negative than would be expected with Kingman's coalescent).

A natural question, then, is how much power these tests have to detect selective sweeps. While a full analysis of this question would require a simulation study, we can obtain some insight from the analytical results presented above. From the values of a_n and b_n in (4.4), which can be found in section 4.1 of Durrett (2002), we see that the standard deviation of the numerator of Tajima's D-statistic is $O(1)$ when the genealogy is given by Kingman's coalescent. However, from the values of c_n and d_n in (5.1), which can be found in section 4.2 of Durrett (2002), we see that the numerator of Fu and Li's D-statistic has a standard deviation which is $O(\log n)$. This means that, for large n, moderate negative values for the numerator of Fu and Li's D-statistic are not strong evidence against the null model of Kingman's coalescent, and thus a test based on Fu and Li's D-statistic will most likely have low power. These observations are consistent with simulation results of Simonsen, Churchill, and Aquadro (1995), who found that Tajima's D-statistic has more power to detect selective sweeps than Fu and Li's D-statistic.

Neither of these tests has the desirable feature of many tests in classical statistics, which is that for all $\alpha > 0$, the power of the level α test tends to 1 as the sample size n tends to

infinity. Indeed, for the problem of detecting recurrent selective sweeps, no such test based on the genealogy of the sample can exist because, with positive probability, none of the selective sweeps affects the genealogy of the n sampled lineages before we get back to the most recent common ancestor. We formulate this observation precisely in the following proposition, which uses the coupling in the proof of Lemma 5.2.

Proposition 5.3. *Let $(\Pi_n(t), t \geq 0)$ be the Λ-coalescent or Ξ-coalescent defined in the proof of Proposition 5.1, and assume that*

$$\sum_{b=2}^{\infty} \frac{\alpha_b}{b^2} < \infty, \tag{5.9}$$

which is slightly weaker than (4.5). Let $(\Upsilon_n(t), t \geq 0)$ be Kingman's coalescent, coupled with $(\Pi_n(t), t \geq 0)$ as in the proof of Lemma 5.2. Then there exists a constant $C > 0$ such that for all n, we have $P(\Upsilon_n(t) = \Pi_n(t) \text{ for all } t) \geq C$.

Proof. Let $T_b = \inf\{t : |\Upsilon_n(t)| = b\}$. Conditional on $\Pi_n(T_b) = \Upsilon_n(T_b)$, the probability that $\Pi_n(t) \neq \Upsilon_n(t)$ for some $t \in [T_b, T_{b-1}]$ is α_b/λ_b. It follows that

$$P(\Upsilon_n(t) = \Pi_n(t) \text{ for all } t) = \prod_{b=2}^{n} \left(1 - \frac{\alpha_b}{\lambda_b}\right).$$

Note that $\alpha_b/\lambda_b \leq 2\alpha_b/[b(b-1)]$ for all b. By (5.9), there exists a positive integer N such that $6\alpha_b/[b(b-1)] \leq 1$ for all $b \geq N$, and if $0 \leq x \leq 1$ then $1 - x/3 \geq e^{-x}$. Putting these results together, we get

$$\begin{aligned} P(\Upsilon_n(t) = \Pi_n(t) \text{ for all } t) &\geq \prod_{b=2}^{N-1} \left(1 - \frac{\alpha_b}{\lambda_b}\right) \prod_{b=N}^{\infty} \exp\left(-\frac{6\alpha_b}{b(b-1)}\right) \\ &= \prod_{b=2}^{N-1} \left(1 - \frac{\alpha_b}{\lambda_b}\right) \exp\left(-\sum_{b=N}^{\infty} \frac{6\alpha_b}{b(b-1)}\right) \geq C, \end{aligned}$$

where the last inequality uses (5.9) again. □

6 Proofs of convergence theorems

In this section, we prove Theorem 2.2 and Proposition 3.2. The proofs use Propositions 2.1 and 3.1 in combination with the Poisson process construction of coalescents with multiple or simultaneous multiple collisions.

Recall the model presented in subsection 2.1 of how the population behaves following a single beneficial mutation. As in subsection 2.1, assume for now that a beneficial mutation occurs at time 0. Let $X(t)$ be the number of chromosomes with the favorable B allele at time t, and let $\tau = \inf\{t : X(t) \in \{0, 2N\}\}$. Let $0 = \xi_0 < \xi_1 < \xi_2 < \dots$ be the times of the proposed replacements, which occur at times of a rate $2N$ Poisson process. Let $0 = \xi'_0 < \xi'_1 < \xi'_2 < \dots$ be the subset of these times at which the number of individuals with the favorable allele changes. As observed in Schweinsberg and Durrett (2004), if $1 \leq k \leq 2N-1$, then $P(X(\xi'_{i+1}) = k+1 | X(\xi'_i) = k) = 1/(2-s)$ and $P(X(\xi'_{i+1}) = k-1 | X(\xi'_i) = k) = (1-s)/(2-s)$. Thus, the number of

chromosomes with the B allele behaves like an asymmetric random walk until it reaches 0 or $2N$. For integers i, j, and k such that $0 \le i \le k \le j \le 2N$ and $i < j$, define

$$p(i,j,k) = P(\inf\{s \ge t : X(s) = j\} < \inf\{s \ge t : X(s) = i\} | X(t) = k),$$

which is the probability that if at some time there are k chromosomes with B, the number of B's will reach j before i. Using the fact that $(1-s)^{\xi'_n}$ is a martingale and applying the Optional Sampling Theorem, we get (see also Durrett (2002) or Lemma 3.1 of Schweinsberg and Durrett (2004))

$$p(i,j,k) = \frac{1-(1-s)^{k-i}}{1-(1-s)^{j-i}}.$$

Therefore, the probability that the beneficial mutation leads to a selective sweep is $p(0,2N,1) = s/(1-(1-s)^{2N})$.

Lemma 6.1 below shows that the length of time that the beneficial allele is present in the population is only $O(\log N)$. Since we speed up time by a factor of N to define the ancestral process, it will follow that for large populations, on the time scale of the ancestral process the lineages that coalesce as a result of a selective sweep coalesce almost at the same time. It is well-known (see Durrett (2002)) that a selective sweep takes time approximately $(2/s)\log(2N)$. However, since a beneficial mutation leads to a selective sweep with probability approximately s, we get a bound on $E[\tau]$ that does not depend on s.

Lemma 6.1. *We have* $E[\tau] \le 4(\log N + 1)$.

Proof. For $1 \le k \le 2N-1$, let $S_k = \#\{i \ge 0 : X(\xi'_i) = k\}$ and $T_k = \#\{i \ge 0 : X(\xi_i) = k\}$, where $\#S$ denotes the cardinality of a set S. Let $q_k = P(X(\xi'_j) \ne k \text{ for all } j > i | X(\xi'_i) = k)$ be the probability that the asymmetric random walk never returns to k. Note that $E[S_k | S_k \ge 1] = 1/q_k$. Also, $P(X(\xi_i) = k \text{ for some } k) = p(0,k,1) = s/(1-(1-s)^k)$. Therefore,

$$E[S_k] = P(S_k \ge 1)E[S_k | S_k \ge 1] = \frac{s}{q_k(1-(1-s)^k)}. \tag{6.1}$$

We have, for $1 \le k \le 2N-1$,

$$\begin{aligned}
q_k &= \left(\frac{1-s}{2-s}\right)[1-p(0,k,k-1)] + \left(\frac{1}{2-s}\right)p(k,2N,k+1) \\
&= \left(\frac{1-s}{2-s}\right)\left[1 - \frac{1-(1-s)^{k-1}}{1-(1-s)^k}\right] + \left(\frac{1}{2-s}\right)\frac{1-(1-s)}{1-(1-s)^{2N-k}} \\
&= \left(\frac{1-s}{2-s}\right)\frac{s(1-s)^{k-1}}{1-(1-s)^k} + \left(\frac{1}{2-s}\right)\frac{s}{1-(1-s)^{2N-k}} \\
&\ge \frac{s}{2-s}\left(\frac{(1-s)^k}{1-(1-s)^k} + 1\right) = \frac{s}{(2-s)(1-(1-s)^k)}.
\end{aligned}$$

It follows from this result and (6.1) that $E[S_k] \le 2-s$ for all k. Schweinsberg and Durrett (2004) calculated that $P(X(\xi_{i+1}) \ne X(\xi_i) | X(\xi_i) = k) = k(2N-k)(2-s)/(2N)^2$. It follows that

$$E[T_k] = E[S_k]\left(\frac{(2N)^2}{k(2N-k)(2-s)}\right) \le \frac{4N^2}{k(2N-k)}.$$

Since $E[\xi_{i+1} - \xi_i] = 1/2N$ for all i, we have

$$E[\tau] = \frac{1}{2N} \sum_{k=1}^{2N-1} E[T_k] \leq \sum_{k=1}^{2N-1} \frac{2N}{k(2N-k)} \leq 2 \sum_{k=1}^{N} \frac{2}{k} \leq 4(\log N + 1),$$

as claimed. □

We now use this result to prove part 2 of Proposition 2.1, which shows that beneficial mutations do not cause lineages to coalesce when the beneficial gene dies out.

Proof of part 2 of Proposition 2.1. Suppose $X(\tau) = 0$ and $\Theta \neq \kappa_0$. Then it can not be true that for all $t \in [0, \tau]$, the n individuals sampled at time τ all have distinct ancestors with the b-chromosome at time t. Therefore, there is an integer i with $\xi_i \leq \tau$ such that one of the following is true:

1. The ancestor at time ξ_i of one of the individuals sampled at time τ has the b allele, but the ancestor of the same individual at time ξ_{i-1} has the B allele because of recombination.

2. There are two individuals in the sample at time τ that have distinct ancestors with the b allele at time ξ_i, but both of them have the same ancestor at time ξ_{i-1}.

We now calculate the probability of these events conditional on $X(\xi_i) = k$, where $1 \leq k \leq N^{1/2}$. We assume $N \geq 2$. For a randomly chosen b chromosome at time ξ_i to have a B chromosome as its ancestor at time ξ_{i-1}, the chosen b chromosome must be the new one born at time ξ_i (which has probability at most $1/(2N-k)$ because $2N-k$ chromosomes have the b allele at time ξ_i), there must be recombination at this time (which happens with probability r), and the ancestor at the site of interest must be a B chromosome (which happens with probability at most $(k+1)/2N$ because $X(\xi_{i-1}) \leq k+1$). Therefore, the probability that all three events occur is at most $r(k+1)/[(2N-k)(2N)] \leq r/N^{3/2}$. Also, at most one pair of b chromosomes at time ξ_i can have the same ancestor at time ξ_{i-1}, so the probability that two randomly chosen b chromosomes coalesce at this time is at most $\binom{2N-k}{2}^{-1} = 2/[(2N-k)(2N-k-1)] \leq 2/N^2$.

By Lemma 6.1, if M is the integer such that $\xi_M = \tau$, then $E[M] \leq (2N)[4(\log N + 1)] = 8N(\log N + 1)$. Since there are n individuals and $\binom{n}{2}$ pairs in the sample, combining these bounds gives

$$P(X(\tau) = 0, X(t) \leq N^{1/2} \text{ for all } t, \text{ and } \Theta \neq \kappa_0) \leq 8N(\log N + 1)\left(\frac{nr}{N^{3/2}} + \frac{n(n-1)}{N^2}\right). \quad (6.2)$$

Note that for $1 \leq k \leq 2N-1$, we have

$$\begin{aligned} P(X(\tau) = 0 \text{ and } X(t) = k \text{ for some } t) &\leq P(X(\tau) = 0 | X(t) = k \text{ for some } t) \\ &= 1 - p(0, 2N, k) = 1 - \frac{1-(1-s)^k}{1-(1-s)^{2N}} \leq (1-s)^k. \end{aligned}$$

Therefore,

$$P(X(\tau) = 0 \text{ and } X(t) > N^{1/2} \text{ for some } t) \leq (1-s)^{N^{1/2}}. \quad (6.3)$$

Combining (6.2) and (6.3), we get

$$P(X_\tau = 0 \text{ and } \Theta \neq \kappa_0) \leq (1-s)^{N^{1/2}} + 8N(\log N + 1)\left(\frac{nr}{N^{3/2}} + \frac{n(n-1)}{N^2}\right).$$

Part 2 of Proposition 2.1 follows because $r \leq C' \log(2N)$ and s is fixed. □

We now consider our model of recurrent selective sweeps and work towards the proof of Theorem 2.2. We will first define a coalescent with multiple collisions. We will then show that this process can be coupled with the ancestral process $(\Psi_N(t), t \geq 0)$ such that, given a finite number of times $0 < u_1 < \cdots < u_m$, the processes agree at these times with high probability.

Recall that K_N is a Poisson point process on $\mathbb{R} \times [-L, L] \times [0, 1]$ with intensity $\lambda \times \mu_N$. We can define another Poisson point process K_N^* on $[0, \infty) \times [-L, L] \times [0, 1]$ which consists of all the points $(-t/N, x, s)$ such that (t, x, s) is a point of K_N and $t \leq 0$. By the Mapping Theorem for Poisson processes (see section 2.3 of Kingman (1993)), K_N^* is a Poisson process with intensity measure $\lambda \times N\mu_N$. The points in K_N^* can be ordered by their first coordinate, so we can write the points as (t_i, x_i, s_i) for positive integers i, where $0 < t_1 < t_2 < \dots$ a.s. Also, define $t_0 = 0$.

We now define a $\mathcal{P}_n$-valued coalescent process $\Pi_N = (\Pi_N(t), t \geq 0)$. Let $\Pi_N(0)$ be the partition κ_0 of $\{1, \dots, n\}$ into singletons. Given $\Pi_N(t_i)$ for some $i \geq 0$, we define $\Pi_N(t)$ for $t_i < t \leq t_{i+1}$ in two steps. First, we let the process obey the law of Kingman's coalescent over the interval (t_i, t_{i+1}), meaning that each possible transition that involves the merging of two blocks happens at rate one. Second, let π_{i+1} be a random partition of $\{1, \dots, n\}$, independent of $(\Pi_N(t), 0 \leq t < t_{i+1})$, such that for an event A_{i+1} of probability s_{i+1}, we have $\pi_{i+1} = \kappa_0$ on A_{i+1}^c and the conditional distribution of π_{i+1} given A_{i+1} is $Q_{p,n}$, where $p = e^{-r_N(x_{i+1}) \log(2N)/s_{i+1}}$. We then define $\Pi_N(t_{i+1})$ to be the coagulation of $\Pi_N(t_{i+1}-)$ by π_{i+1}.

The lemma below states that the coalescent process Π_N that we have just defined is a coalescent with multiple collisions.

Lemma 6.2. *Let η_N be the measure on $(0, 1]$ such that*

$$\eta_N([y, 1]) = \int_{-L}^{L} \int_0^1 s 1_{\{e^{-r_N(x)\log(2N)/s} \geq y\}} \, N\mu_N(dx \times ds)$$

for all $y \in (0, 1]$. Let $\Lambda_{0,N}$ be the measure on $(0, 1]$ such that $\Lambda_{0,N}(dx) = x^2 \eta_N(dx)$, and let $\Lambda_N = \delta_0 + \Lambda_{0,N}$. Then the process $(\Pi_N(t), t \geq 0)$ is the $\mathcal{P}_n$-valued Λ_N-coalescent.

Proof. Let K_N' be the point process on $[0, \infty) \times \mathcal{P}_n$ consisting of the points (t_i, π_i). By the Marking Theorem for Poisson processes (see section 5.2 of Kingman (1993)), K_N' is also a Poisson point process. Given (t_i, x_i, s_i), the partition π_i has distribution $Q_{p,n}$, where $p = e^{-r_N(x_i)\log(2N)/s_i}$, conditional on an event of probability s_i and otherwise is κ_0. Therefore, the intensity measure of K_N' is given by $\lambda \times H$, where, for $\pi \neq \kappa_0$,

$$\begin{aligned} H(\pi) &= \int_{-L}^{L} \int_0^1 s Q_{e^{-r_N(x)\log(2N)/s},n}(\pi) \, N\mu_N(dx \times ds) \\ &= \int_0^1 Q_{p,n}(\pi) \, \eta_N(dp) = \int_0^1 Q_{p,n}(\pi) p^{-2} \Lambda_{0,N}(dp). \end{aligned}$$

By comparing this with (3.3) and recalling that Π_N follows the law of Kingman's coalescent over the intervals (t_{i-1}, t_i), we conclude that Π_N is the Λ_N-coalescent. □

The next lemma states that it is unlikely for there to be a beneficial allele in the population at any fixed time. Recall that $\mathcal{T}_N = \{t : (t, x, s) \text{ is a point in } K_N \text{ for some } x \text{ and } s\}$.

Lemma 6.3. *There exists a constant C, not depending on N, such that for any fixed $y \in \mathbb{R}$, we have $P(y \in [t, \tau(t)] \text{ for some } t \in \mathcal{T}_N) \le (C \log N)/N$.*

Proof. The points of $\mathcal{T}_N$ form a Poisson process on $\mathbb{R}$ of rate γ_N, where $\gamma_N = \mu_N([-L, L] \times [0, 1])$. Recall from Lemma 6.1 that if τ denotes the amount of time for which a beneficial allele is present in between 1 and $2N - 1$ members of the population, then $E[\tau] \le 4(\log N + 1)$. Therefore,

$$\begin{aligned} P(y \in [t, \tau(t)] \text{ for some } t \in \mathcal{T}_N) &\le \int_{-\infty}^{y} P(\tau \ge y - x)\gamma_N \, dx \\ &= \gamma_N \int_0^{\infty} P(\tau \ge x) \, dx = \gamma_N E[\tau] \le 4\gamma_N(\log N + 1). \end{aligned}$$

Since the measures $N\mu_N$ converge weakly to μ, the sequence $(N\gamma_N)_{N=1}^{\infty}$ converges to $\mu([-L, L] \times [0, 1])$ and therefore is bounded. The lemma follows. □

We now show how to couple the processes Ψ_N and Π_N so that they agree at a given finite set of times with high probability. We first consider how the ancestral process Ψ_N behaves around the times $t_1, t_2, \ldots$. For positive integers i, let $\tau_i = -\tau(-Nt_i)/N$. We have $-Nt_i \in \mathcal{T}_N$. However, recall from subsection 2.2 that not all points in $\mathcal{T}_N$ are in $\mathcal{T}'_N$ because some potential mutations are discarded to avoid overlapping selective sweeps. When $-Nt_i \in \mathcal{T}'_N$, there is a beneficial allele in the population during the time interval $[-Nt_i, \tau(-Nt_i))$, and this affects the process Ψ_N over the interval $[\tau_i, t_i]$.

For each i such that $-Nt_i \in \mathcal{T}'_N$, we can define a random partition $\theta_i \in \mathcal{P}_n$ by choosing n individuals from the population at time $\tau(-Nt_i)$ and declaring two integers j and k to be in the same block of θ_i if and only if the jth and kth individuals chosen got their allele at the neutral site of interest from the same ancestor at time $-Nt_i$. If $\tau_i > 0$ and the partition $\Psi_N(\tau_i)$ contains b_i blocks, we can choose the n individuals at time $\tau(-Nt_i)$ by first picking the b_i individuals that are ancestors of the n individuals that were sampled at time zero, and then choosing the remaining $n - b_i$ at random. This will ensure that, for i such that $-Nt_i \in \mathcal{T}'_N$ and $\tau_i > 0$, the random partition $\Psi_N(t_i)$ is the coagulation of $\Psi_N(\tau_i)$ by θ_i.

Moreover, the conditional distribution of θ_i given (t_i, x_i, s_i) and given that $-Nt_i \in \mathcal{T}'_N$ is the same as the distribution of the random partition Θ defined in subsection 2.1, when the selective advantage is s_i and the recombination probability is $r_N(x_i)$. Recall that when a beneficial mutation occurs in the population with selective advantage s_i, it spreads to the entire population with probability $s_i/(1 - (1 - s_i)^{2N})$. Therefore, by Proposition 2.1, the distribution of Θ is approximately that of a random partition that has distribution $Q_{p,n}$, where $p = e^{-r_N(x_i)\log(2N)/s_i}$, on an event of probability $s_i/(1-(1-s_i)^{2N})$ and is κ_0 on the complementary event. However, this is the same as the conditional distribution of π_i given (t_i, x_i, s_i), except we have $s_i/(1-(1-s_i)^{2N})$ instead of s_i. It thus follows from Proposition 2.1 that we can couple the partitions θ_i and π_i such that for any $\delta > 0$,

$$P(\theta_i \ne \pi_i \text{ and } -Nt_i \in \mathcal{T}'_N | (t_i, x_i, s_i)) \le \frac{C_\delta}{\log N} + 1_{\{s_i < \delta\}}, \tag{6.4}$$

where C_δ is a constant that depends on δ. Note that we only get the $O(1/(\log N))$ bound when $s_i \ge \delta$ because of the assumption in Proposition 2.1 that s is fixed.

Finally, we consider the processes during the intervals (t_i, t_{i+1}). The process Π_N behaves like Kingman's coalescent during these intervals. Let

$$\mathcal{I}_N^* = \bigcup_{i=1}^{\infty} [\tau_i, t_i].$$

The process Ψ_N behaves like Kingman's coalescent during the intervals in $(0,\infty)\backslash\mathcal{I}_N^*$ because the population follows the Moran model during the corresponding intervals. Therefore, if $\Pi_N(t_i) = \Psi_N(t_i)$, we can couple the processes so that $\Pi_N(t) = \Psi_N(t)$ for all $t \in [t_i, \phi_i)$, where $\phi_i = \inf\{t > t_i : t \in \mathcal{I}_N^*\}$.

Proposition 6.4. *Suppose the processes Π_N and Ψ_N are coupled in the manner described above. Let $0 < u_1 < \cdots < u_m$ be fixed times. Let $\epsilon > 0$. For sufficiently large N, we have*

$$P(\Pi_N(u_i) \neq \Psi_N(u_i) \text{ for some } i \in \{1, \ldots, m\}) < \epsilon. \tag{6.5}$$

Proof. Let $K = \sup\{k : t_k \leq u_m\}$. Suppose the following conditions hold:

1. For $i = 1, \ldots, m$, we have $u_i \notin \mathcal{I}_N^*$.
2. For all positive integers i, we have $\tau_i > 0$.
3. For $i = 1, \ldots, K$, we have $-Nt_i \in \mathcal{T}_N'$.
4. For $i = 1, \ldots, K$, we have $\Pi_N(\tau_i) = \Pi_N(t_i-)$.
5. For $i = 1, \ldots, K$, we have $\theta_i = \pi_i$.

Conditions 2 and 3 imply that

$$0 = t_0 < \tau_1 < t_1 < \tau_2 < t_2 < \cdots < \tau_K < t_K \leq u_m.$$

Condition 1 with $i = m$ implies further that $\tau_j > u_m$ for all $j > K$, so $(t_K, u_m] \subset \mathbb{R}\backslash\mathcal{I}_N^*$. We know that $\Pi_N(t_0) = \Psi_N(t_0) = \kappa_0$. Suppose, for some $i \in \{0, \ldots, K-1\}$, that $\Pi_N(t_i) = \Psi_N(t_i)$. Then the coupling gives $\Pi_N(t) = \Psi_N(t)$ for all $t \in [t_i, \tau_{i+1})$. Condition 4 gives $\Pi_N(\tau_{i+1}) = \Pi_N(t_{i+1}-)$. Conditions 2 and 3 imply that $\Psi_N(t_{i+1})$ is the coagulation of $\Psi_N(\tau_{i+1})$ by θ_{i+1}. Since $\Pi_N(t_{i+1})$ is the coagulation of $\Pi_N(t_{i+1}-)$ by π_{i+1}, condition 5 ensures that $\Pi_N(t_{i+1}) = \Psi_N(t_{i+1})$. Thus, $\Pi_N(t_i) = \Psi_N(t_i)$ for $i = 0, 1, \ldots, K$, and the coupling combined with the fact that $\Pi_N(t_K) = \Psi_N(t_K)$ gives $\Pi_N(t) = \Psi_N(t)$ for all $t \in (t_K, u_m]$. Thus, we have $\Pi_N(t) = \Psi_N(t)$ for all $t \in [0, u_m] \setminus \mathcal{I}_N^*$. Therefore, by condition 1, $\Pi_N(u_i) = \Psi_N(u_i)$ for $i = 1, \ldots, m$. It thus remains only to show that conditions 1 through 5 occur with high probability. For the rest of the proof, we allow the constant C to change from line to line.

If $u_i \in \mathcal{I}_N^*$, then there exists $t \in \mathcal{T}_N$ such that $-Nu_i \in [t, \tau(t)]$. Therefore, by Lemma 6.3,

$$P(u_i \in \mathcal{I}_N^* \text{ for some } i \in \{1, \ldots, m\}) \leq \frac{C \log N}{N}.$$

Likewise, if $\tau_i < 0$ for some i, then $-Nt_i \leq 0 \leq \tau(-Nt_i)$ and $-Nt_i \in \mathcal{T}_N$. It follows that $P(\tau_i < 0 \text{ for some } i) \leq C(\log N)/N$ by Lemma 6.3.

To deal with conditions 3, 4, and 5, let $l_j = ju_m/N$ for $j = 0, 1, \ldots, N$, and define the intervals $I_1, \ldots, I_N$ by $I_j = [l_{j-1}, l_j]$. Note that the number of the points t_i in an interval I_j is Poisson with mean $u_m\gamma_N$. Therefore, the probability that some point t_i falls in I_j is at most $u_m\gamma_N \le C/N$. The probability that two or more points fall in I_j is at most $u_m^2\gamma_N^2 \le C/N^2$. If there is a point $t_i \in I_j$ with $-Nt_i \notin \mathcal{T}'_N$, then either there are two points in I_j or there is one point in I_j and $l_j \in \mathcal{I}^*_N$. The event that there is at least one point in I_j is independent of the event that $l_j \in \mathcal{I}^*_N$, so using Lemma 6.3 again, the probability that both occur is at most $C(\log N)/N^2$.

When the number of blocks in the coalescent is at most n, the total transition rate of the process Π_N is bounded by $\binom{n}{2} + N\gamma_N$. The probability that there is any point t_i in I_j is at most C/N, so by Lemma 6.1,

$$P(\Pi_N(\tau_i) \neq \Pi_N(t_i-) \text{ for some } t_i \in I_j) \le \frac{C}{N}\left(\binom{n}{2} + N\gamma_N\right)E[t_i - \tau_i] \le \frac{C\log N}{N^2}.$$

Finally, we may choose δ small enough that $P(s_i \le \delta) < \epsilon$, and then (6.4) gives

$$P(\theta_i \neq \pi_i | t_i \in I_j) < \epsilon + \frac{C_\delta}{\log N},$$

where C_δ is a constant that depends on δ. Therefore, $P(\theta_i \neq \pi_i \text{ for some } t_i \in I_j) \le \epsilon/N + C_\delta/(N\log N)$. Since there are only N intervals I_j, we can add these bounds to show that the probability that conditions 1 through 5 all hold is at least $1 - \epsilon$ for sufficiently large N, which implies the statement of the proposition. □

Proof of Theorem 2.2. Let $0 < u_1 < \cdots < u_m$ be fixed times. Let $\epsilon > 0$. Define Λ_N as in Lemma 6.2, and let Π_N be a $\mathcal{P}_n$-valued Λ_N-coalescent. In view of Proposition 6.4, it suffices to show that for all $\pi_1, \ldots, \pi_m \in \mathcal{P}_n$, we have

$$|P(\Pi_N(u_i) = \pi_i \text{ for all } i \in \{1, \ldots, m\}) - P(\Pi(u_i) = \pi_i \text{ for all } i \in \{1, \ldots, m\})| < \epsilon$$

for sufficiently large N. Therefore (see Pitman (1999)), it suffices to show that the measures Λ_N converge weakly to Λ. Thus, we need to show (see Billingsley (1999), Theorem 2.1) that for any bounded uniformly continuous function h on $[0, 1]$, we have $\int_0^1 h(x)\,\Lambda_N(dx) \to \int_0^1 h(x)\,\Lambda(dx)$ as $N \to \infty$. By the definitions of Λ_N and Λ, it suffices to show that $\int_0^1 h(x)\,\eta_N(dx) \to \int_0^1 h(x)\,\eta(dx)$ as $N \to \infty$ for any bounded uniformly continuous function h on $(0, 1]$. By the definitions of the measures η_N and η, this is equivalent to showing that

$$\lim_{N\to\infty} \int_{-L}^{L}\int_0^1 sh\big(e^{-r_N(x)\log(2N)/s}\big)\,N\mu_N(dx \times ds) = \int_{-L}^{L}\int_0^1 sh\big(e^{-r(x)/s}\big)\,\mu(dx \times ds) \qquad (6.6)$$

for any bounded uniformly continuous function h on $(0, 1]$. However, it is easy to deduce (6.6) from the boundedness and uniform continuity of h, the uniform convergence of $(\log 2N)r_N$ to r, the continuity of r, and the weak convergence of the measures $N\mu_N$ to μ. □

Proof of Proposition 3.2. Proposition 3.2 can be proved by repeating the proof of Proposition 6.4 with minor changes. To prove the first part of the proposition, we construct the coalescent process Π_N as before. Because $N\mu_N = \mu$ and $\log(2N)r_N = r$ for all N, we have $\Lambda_N = \Lambda$ for all N. It follows from Lemma 6.2 that Π_N is a Λ-coalescent for all N. Thus, it suffices to show

(6.5), but with $C/(\log N)$ on the right-hand side instead of ϵ. Because we are assuming that μ is concentrated on $[-L, L] \times [\epsilon, 1]$ for some $\epsilon > 0$, we can choose $\delta = \epsilon$ and drop the indicator from the right-hand side of (6.4) to get a bound of $C_\epsilon/(\log N)$. We then obtain $C/(\log N)$ on the right-hand side of (6.5) by following the same steps as before.

To prove the second part of Proposition 3.2, we modify the definition of Π_N. Conditional on A_i, we give π_i the distribution $Q_{R(r_N(x_i)/s_i,\lfloor 2Ns_i\rfloor),n}$. We set $\pi_i = \kappa_0$ on A_i^c. The intensity measure of K'_N is then given by $\lambda \times J$, where, for all $\pi \neq \kappa_0$, we have

$$\begin{aligned} J(\pi) &= \int_{-L}^{L} \int_0^1 s Q_{R(r_N(x)/s,\lfloor 2Ns\rfloor),n}(\pi)\, N\mu_N(dx \times ds) \\ &= \int_\Delta Q_{\delta_x,n}(\pi)\, G_N(dx) = \int_\Delta Q_{\delta_x,n}(\pi) \Big(\sum_{j=1}^{\infty} x_j^2 \Big)^{-1} \Xi_{N,0}(dx). \end{aligned}$$

By comparing this with (3.2), we see that the process Π_N is a Ξ_N-coalescent. It follows from Proposition 3.1 that we can replace $C_\delta/(\log N)$ on the right-hand side of (6.4) by $C_\delta/(\log N)^2$. This gives the second part of the proposition. □

References

N. H. Barton (1995). Linkage and the limits to natural selection. *Genetics*, **140**, 821-841.

N. H. Barton (1998). The effect of hitch-hiking on neutral genealogies. *Genet. Res.* **72**, 123-133.

N. H. Barton, A. M. Etheridge, and A. K. Sturm (2004). Coalescence in a random background. *Ann. Appl. Probab.* **14**, 754-785.

J. Bertoin and J.-F. LeGall (2003). Stochastic flows associated to coalescent processes. *Probab. Theory Relat. Fields*, **126**, 261-288.

P. Billingsley (1999). *Convergence of Probability Measures.* 2nd ed. New York, Wiley.

J. M. Braverman, R. R. Hudson, N. L. Kaplan, C. H. Langley, and W. Stephan (1995). The hitchhiking effect on the site frequency spectrum of DNA polymorphisms. *Genetics*, **140**, 783-796.

P. Donnelly and T. G. Kurtz (1999). Genealogical processes for Fleming-Viot models with selection and recombination. *Ann. Appl. Prob.* **9**, 1091-1148.

R. Durrett (2002). *Probability Models for DNA Sequence Evolution.* New York, Springer-Verlag.

R. Durrett and J. Schweinsberg (2004a). Approximating selective sweeps. *Theor. Popul. Biol.* **66**, 129-138.

R. Durrett and J. Schweinsberg (2004b). Power laws for family sizes in a duplication model. Preprint. Available at http://front.math.ucdavis.edu/math.PR/0406216.

Y. X. Fu and W. H. Li (1993). Statistical tests of neutrality of mutations. *Genetics*, **133**, 693-709.

P. J. Gerrish and R. E. Lenski (1998). The fate of competing beneficial mutations in an asexual population. *Genetica*, **102/103**, 127-144.

J. H. Gillespie (2000). Genetic drift in an infinite population: the pseudohitchhiking model. *Genetics*, **155**, 909-919.

N. L. Kaplan, R. R. Hudson, and C. H. Langley (1989). The "hitchhiking effect" revisited. *Genetics*, **123**, 887-899.

Y. Kim and W. Stephan (2003). Selective sweeps in the presence of interference among partially linked loci. *Genetics*, **164**, 389-398.

J. F. C. Kingman (1978). The representation of partition structures. *J. London Math. Soc.* **18**, 374-380.

J. F. C. Kingman (1982). The coalescent. *Stochastic Process. Appl.* **13**, 235-248.

J. F. C. Kingman (1993). *Poisson processes*. Oxford, Clarendon Press.

S. M. Krone and C. Neuhauser (1997). Ancestral processes with selection. *Theor. Popul. Biol.* **51**, 210-237.

J. Maynard Smith and J. Haigh (1974). The hitchhiking effect of a favorable gene. *Genet. Res.* **23**, 23-35.

M. Möhle and S. Sagitov (2001). A classification of coalescent processes for haploid exchangeable population models. *Ann. Probab.* **29**, 1547-1562.

P. A. P. Moran (1958). Random processes in genetics. *Proc. Cambridge Philos. Soc.* **54**, 60-71.

C. Neuhauser and S. M. Krone (1997). The genealogy of samples in models with selection. *Genetics*, **145**, 519-534.

J. Pitman (1999). Coalescents with multiple collisions. *Ann. Probab.* **27**, 1870-1902.

M. Przeworski (2002). The signature of positive selection at randomly chosen loci. *Genetics*. **160**, 1179-1189.

S. Sagitov (1999). The general coalescent with asynchronous mergers of ancestral lines. *J. Appl. Probab.* **36**, 1116-1125.

S. Sagitov (2003). Convergence to the coalescent with simultaneous multiple mergers. *J. Appl. Probab.* **40**, 839-854.

J. Schweinsberg (2000). Coalescents with simultaneous multiple collisions. *Electron. J. Probab.* **5**, 1-50.

J. Schweinsberg (2003). Coalescent processes obtained from supercritical Galton-Watson processes. *Stochastic Process. Appl.* **106**, 107-139.

J. Schweinsberg and R. Durrett (2004). Random partitions approximating the coalescence of lineages during a selective sweep. Preprint. Available at http://front.math.ucdavis.edu/math.PR/0411069.

K. Simonsen, G. A. Churchill, and C. F. Aquadro (1995). Properties of statistical tests of neutrality for DNA polymorphism data. *Genetics*, **141**, 413-429.

W. Stephan, T. Wiehe, and M. W. Lenz (1992). The effect of strongly selected substitutions on neutral polymorphism: Analytical results based on diffusion theory. *Theor. Popul. Biol.* **41**, 237-254.

F. Tajima (1989). Statistical method for testing the neutral mutation hypothesis by DNA polymorphism. *Genetics*, **123**, 585-595.